

Survey of Multiscale and Multiphysics Applications and Communities

Derek Groen, Stefan J. Zasada, Peter V. Coveney

Centre for Computational Science

University College London

London, United Kingdom

E-mail: d.groen@ucl.ac.uk, p.v.coveney@ucl.ac.uk

Abstract—Multiscale and multiphysics applications are becoming increasingly common, and many researchers focus on combining existing models to construct combined multiscale models. Here we present a concise review of multiscale applications and their source communities. We assess and compare the methods they use to construct their multiscale models and we characterize areas where inter-disciplinary multiscale collaboration could be particularly beneficial. We conclude that multiscale computing has become increasingly popular in recent years, that different communities adopt very different organizational approaches, and that simulations on a length scale of a few metres and a time scale of a few hours can be found in many of the multiscale research domains. Sharing multiscale methods specifically geared towards these scales between communities may therefore be particularly beneficial.

Keywords—multiscale computing; application review; multi-scale software; multiscale communities

I. INTRODUCTION

Many physical problems we seek to understand nowadays are complex in nature, and consist of separate physical processes that each contribute to the problem as a whole. These processes, for example in biology the interactions between binding molecules or the dynamics occurring at the cellular level, each take place on a specific space scale or time scale. For example, the interactions between molecules typically take place on a space scale of several nanometers and a time scale of a number of nanoseconds. However, to model the interactions on the cellular level will require a considerably larger space and time scale. Historically many problems were investigated by modeling or simulating a physical process in isolation, and from the outcome of that exercise, determine its contribution to the overall (complex) physical problem. In the last two decades a new approach has become increasingly widespread, where researchers construct models and simulations that capture multiple physical processes, each of which is represented as a *submodel*, operates on a different space or time scale and has the potential to influence other processes. This approach is now known as *multiscale modelling* or *multiscale simulation*. Within this article we use the term multiscale modelling to refer to both the multiscale modelling and simulation of physical problems, and the term *multiscale application* to refer to the program used to do the modelling. In turn, we use the term *subcode* to refer to the implementations of each submodel.

A. Multiphysics modelling

When a model captures multiple physical processes, and each of these processes capture a different type of physics, it is commonly referred to as *multiphysics modelling* or *multiphysics simulation*. For example, a model of a star cluster that resolves Newtonian gravitational interactions using one submodel and the aging of stars using another submodel is considered to be a multiphysics submodel, even if these models were (hypothetically) to operate on the same space and time scale. However, a star cluster model that uses two different submodels for the Newtonian gravitational interaction of stars is generally not considered to be multiphysics, even when these models may be applied on a different space or time scale.

Multiscale and multiphysics modelling are therefore two different concepts, but they do have one prime commonality in both cases the models consist of a number of submodels which have been combined (or *coupled*). A major challenge in multiscale as well as multiphysics modelling lies in coupling these submodels such that the overall model is both accurate enough to be scientifically relevant and reproducible, and efficient enough to be executed conveniently by modern compute resources.

B. Multiscale and multiphysics applications

Multiscale and multiphysics applications are present in a wide range of scientific and engineering communities. By its nature, multiscale modeling is highly interdisciplinary, with developments occurring independently across research domains. Here we review a range of multiscale applications and communities that reside within different scientific domains. We describe several major projects for each domain and present the results of our investigation on the popularity of multiscale simulation and modeling. We find that multiscale methods are adopted in hundreds of projects both in the EU and US, and that the popularity of multiscale simulation and modeling has increased considerably in recent years.

We also illustrate approaches to construct multiscale simulations in different scientific domains, and compare some of the characteristics of the multiscale communities in these domains. Additionally, we present a comparison between coupling frameworks, and point out potential areas where

interdisciplinary collaborations would be particularly beneficial. Within this survey we cover major areas of multiscale simulation and modeling activities, but this review is by no means exhaustive. For readability reasons we provide only a limited number of references here. However, a full literature list is available as a web-based supplement for those who wish to delve more deeply into the work performed by the various multiscale simulation and modeling communities.

C. Related work

Aside from numerous publications, project websites and domain-specific reviews, we have identified a few sources which provide information on multiscale simulations in various scientific domains. One such source of information is the *Journal of Multiscale Modeling and Simulation* (epubs.siam.org/mms), which defines itself as an interdisciplinary journal focusing on the fundamental modeling and computational principles underlying various multiscale methods. The *Journal of Multiscale Modeling* (www.worldscinet.com/jmm/) is also targeted at multiscale modeling in general. There are also several books which present multiscale research in a range of domains [1], [2], as well as dozens of multiscale modeling workshops such as the Multiscale Materials Meeting conference (www.mrs.org.sg/mmm2012) or the Distributed Multiscale Computing (www.computationalscience.nl/dmc2011/). In addition, an investigation on the modeling aspects of multiscale simulations, emphasizing simulations using Cellular Automata, is provided in [3]. Many of the model classifications in this work can be reused for multiscale problems in other areas of science.

II. OVERVIEW OF MULTISCALE COMMUNITIES

A. Astrophysics

Astrophysics is a scientific domain which features a multitude of active multiscale projects. Due to the inherently large size and complex nature of many astrophysical problems, incorporating of multiple scales and multiple types of physics allows researchers to more accurately model a wide range of astrophysical phenomena. As a result, researchers developed multiscale codes in fields such as cosmology [4], star cluster dynamics [5], [6], thermonuclear supernovae [7] and modeling space weather [8].

Perhaps the largest and most versatile multiscale framework currently in use in astrophysics is the Astrophysical Multipurpose Software Environment (AMUSE, www.amusecode.org). The aim of AMUSE [9], as a successor of the MUSE toolkit [10], is to provide a software framework for multiscale astrophysical simulations. This framework connects existing (parallel) codes through a layer of Python and MPI to form coupled simulations. There are numerous astrophysical scenarios in which AMUSE has been applied, for example for coupling a gravitational N-body simulation with a stellar evolution code to model both

the dynamical movements and the aging of the stars in a star cluster [10].

The FLASH code [11] couples hydrodynamics with magnetic fields and is applied to simulate the surfaces of compact stars such as white dwarves and neutron stars. The latest version, FLASH 4 [12], consists of inter-operable modules that can be combined to generate different applications. The FLASH architecture allows an arbitrary number of alternative implementations of its components to co-exist and be interchanged with each other, which results in greater flexibility. Furthermore, a simple and elegant mechanism exists for customization of code functionality without the need to modify the core implementation of each component.

B. Biology

Biological systems span many orders of magnitude through the length and time scales in a continuous way, from the smallest molecular scale up to the whole body. The sequence from the genome, proteome, metabolome, physiome to health comprises multi-scale systems biology of the most ambitious kind [13], [14], [15], [16], [17]. Schnell et al. [18] provide a pedagogical introduction to this concept and the field of multi-scale modeling in biology, while Sloot et al. [19] provide an extensive prior review of multiscale modeling in computational biomedicine. In addition, Dada et al. [20] give a general overview of the multiscale modeling projects in biology. The developments in this domain have led to the emergence of several multiscale coupling tools, such as GridSpace (dice.cyfronet.pl/gridspace) and MUSCLE (www.irmb.bau.tu-bs.de/muscle/).

The Virtual Physiological Human (VPH) Initiative is a large and active community in the multiscale computing domain, which has received substantial support throughout the EU 7th framework program ICT-VPH. Multiscale simulations and models are of fundamental importance within the VPH. VPH-like activities are represented in Europe (VPH-NoE, www.vph-noe.eu), USA (IMAG and the Multi-scale Modeling Consortium, www.imagwiki.nibib.nih.gov), Japan (e.g. the group of Himeno at RIKEN, www.riken.jp/engn/r-world/research/lab/rpcs/) and world-wide in terms of the Physiome project [21] (www.physiome.org).

C. Energy

The need for multiscale simulations in the energy domain has been identified by a number of groups (e.g., [22]). Within the EU hundreds of researchers participate in modeling present fusion devices and preparing for modeling ITER (www.iter.org) and its successor the DEMOnstration Power Plant (DEMO). These models are often multiscale in nature, and require coupling dozens of different submodels, each of which resolves different physical aspects on a different spatial scale. Several tools emerged that assist in coupling fusion applications, such as the Universal Access Layer (UAL [23]), the Framework Application for Core-Edge

Transport Simulations (FACETS, www.facetsproject.org) and the Integrated Plasma Simulator (cswim.org/ips/). Much of the effort within Europe centers around the EFDA Task Force on Integrated Tokamak Modeling (www.efda-itm.eu), where for example they use multiscale methods to model the lifetime of fusion reactor materials. Additionally, the developments in the GriPhyN high energy physics computing project (www.griphyn.org) resulted in a generalized toolkit for workflow-style multiscale simulations (Swift [24]).

D. Engineering

The engineering domain comprises numerous multiscale projects. Example efforts aiming at multiscale modeling include the FLOMANIA and DESIDER projects [25], [26] in the aeronautics domain, research on multiscale methods for non-equilibrium physics [27], investigations on catalysis and reaction engineering [28], [29], stochastic simulations of kinetic theory models [30] and the coupling of atomistic and continuum methods in hydrology [31], [32]. Additionally, several universities have set up centres for multiscale modeling and simulation, such as Stanford University (me.stanford.edu/research/multiscale.html) and Carnegie-Mellon University (www.ices.cmu.edu/cm2em/). The *International Journal of Multiscale Computational Engineering* (www.begellhouse.com/journals/61fd1b191cf7e96f) is a new journal dedicated to multiscale simulation. Additionally, a general review of multiscale simulation in science and engineering is provided in [1].

E. Environmental science

Environmental science is intrinsically multiscale and multi-process, ranging across topics such as ecology studies, climate modeling, geosciences and hydrology. There is a diverse collection of multiscale projects within this domain, including hydrology simulation [33], weather forecasting [34], [35], climate modeling [36] and disaster prediction [37]. Klein et al. [38] provide a broad review of multiscale (fluid dynamics) methods in meteorology. Researchers within this domain have also developed several general-purpose toolkits, such as the Model Coupling Toolkit [39], the Pyre framework [40] (www.cacr.caltech.edu/projects/pyre), OpenPALM (www.cerfacs.fr/globc/PALM_WEB), OASIS [41] and OpenMI [42].

The European Network for Earth System Modelling (www.enes.org) is a large consortium focused on developing of a European network for Earth system (multiscale) modeling. The ENES consortium has helped to establish the ENSEMBLES project (ensembles-eu.metoffice.com), where a consortium of 70 partners use multiscale ensemble simulations, varying across temporal and spatial scales, to simulate the Earth system for a range of applications. These include climate predictions, as well as the resolution of physical, chemical, biological and human-related feedback processes.

This Earth system also includes water resource and land use, air quality issues, as well as carbon cycle feedbacks.

F. Materials science

Materials science applications are inherently multiscale. The macroscopic level may be the one in which we can see and interact with materials, but the interactions that define the properties of many materials occur on the microscopic level. Linking our understanding of the physical world at very short scales with the observable behaviour at the macro-scale is a major focus within this area of science, and the applications are extremely varied. The topics covered in these projects range from multiscale modeling of radiation damage (e.g., RADINTERFACES [43]) to modeling of multilayered surface systems (e.g., M3-2S [44]) and the modeling of porous materials [45], multiscale heterogeneous modeling of solids [46]. The book by Attinger and Koumoutsakos comprehensively presents a large number of projects within the materials sciences [2].

Additionally, MMM@HPC (www.multiscale-modelling.eu) is an ongoing EU project which develops infrastructure and software for multiscale materials modelling. An example of distributed multiscale materials modeling is the clay-polymer nanocomposites application within the MAPPER project [47] (www.mapper-project.eu), a project which features multiscale applications from a wide range of domains. We have not found any coupling frameworks within this domain, but there is at least one public tool, VOTCA (www.votca.org), which facilitates the process of converting data from atom-level simulations to a larger and more coarse-grained level. The book by Attinger and Koumoutsakos presents a large number of projects in the materials sciences as well [2].

G. Other domains

One other area which features a large number of multiscale simulations is the scientific domain of fluid dynamics. Active areas of multiscale research include the application of hybrid multiscale algorithms to model complex fluids [48], [46], [49], multi-physics approaches to biofluid simulations [50] and multiscale analysis of magnetorheological fluids [51]. Additionally, both the *International Journal of Multiscale Computational Engineering* [52] and the *Journal of Multiscale Modelling* [53] contain many multiscale research articles within this field.

The multiscale modeling and simulation efforts within fluid dynamics frequently take place within the context of other scientific domains, such as biology in the case of blood flow simulations, and environmental science in the case of oceanic simulations. To accommodate this, we have not sought to treat fluid dynamics as a separate domain, but prefer to categorize the projects in accordance with their application domain.




Figure 1. Examples of a acyclically (left) and two cyclically coupled (middle and right) multiscale models. Submodels are indicated by blue boxes, and data transfers by arrows. On the right we provide a cyclically coupled model where the submodels are executed concurrently. The concurrent execution is frequently managed by a software tool that supports cyclic coupling, which we indicate there with a yellow ellipse.

Overall, the six domains described in this work represent major areas where multiscale simulations are frequently applied. Having performed an extensive search, we did find a number of multiscale projects outside these domains. The vast majority of these projects concern theoretical mathematical modeling of multiscale problems, and only indirectly relate to the other scientific fields in our survey.

III. REVIEW OF MULTISCALE COMMUNITIES

In this section we describe some key characteristics from several multiscale communities. The main purpose of this comparison is to find common aspects between different communities and multiscale coupling tools, and to identify potential areas for systematic and reusable approaches for multiscale simulation. Within this paper we distinguish between two multiscale simulation methods: *acyclically coupled* simulations and *cyclically coupled* simulations. Acyclically coupled simulations are applications where subcodes are run, producing results which in turn are used as input for the execution of a subsequent subcode. The most characteristic aspect of acyclically coupled simulations is that there are no cases where two or more subcodes are mutually dependent of each other during execution. Cyclically-coupled simulations do have this mutual dependency, and require at least some of the subcodes to be either run concurrently or in alternating fashion as they are mutually dependent. We show several schematic examples of multiscale models, both using acyclic coupling and cyclic coupling, in Fig. 1.

We provide an overview of the current state of multiscale computing in six scientific domains in Table I. All these domains feature at least several projects where parallel simulation codes have been applied to construct multiscale simulations. Geographically distributed multiscale simulations are less common, although we did find at least one example for five of the six domains. Distributed multiscale simulations are slightly more common in the biology do-

main, with three disjoint applications being present within the MAPPER project alone.

We provide a schematic graphical overview of the spatial and temporal scales commonly adopted by different communities in Fig. 2. Each research discipline has a unique range of spatial and temporal scales, with for example quantum mechanical simulations in the materials sciences operating on the smallest scales and cosmological dark matter simulations on the largest scales. The time and space scale range of each discipline is given respectively by the visually observed height and width of the corresponding parallelograms. Here, relatively small parallelograms (as is the case for mechanical engineering and environmental science) point to a higher probability of overlapping scales between subcodes within those disciplines. When spatial and temporal scales between subcodes overlap, cyclic interactions between submodels are essential to obtain an accurate result, and it becomes difficult to accurately model the system using acyclic coupling alone. More information on the consequences of overlapping spatial and temporal scales can be found in Hoekstra et al. [3].

When we compare different disciplines, we observe a roughly linear trend between the temporal scale and the spatial scale of multiscale simulations. This correlation is to be expected as shorter-range interactions tend to operate on shorter time scales as well. We also observe that phenomena within a spatial range between 10^{-3}m and 10^3m and a temporal range between 10^1s and 10^5s are commonly addressed in a large number of scientific disciplines.

This region of overlap may be of particular interest when opting for interdisciplinary approaches or reusable multiscale simulation tools. Additionally, when a particularly high accuracy is required within a simulation operating on these overlapping scales, it may become increasingly important to incorporate phenomena from other overlapping scientific disciplines, given that these phenomena are sufficiently proximate.

A. Prevalence of multiscale research

To gain some understanding of the size of existing multiscale research communities we have explored several project data bases from large funding agencies. These include the European Community Research and Development Information Service (CORDIS), as well as the project databases of the National Institute for Health (NIH), the Department of Energy (DOE) and the US National Science Foundation (NSF). We found the projects by first selecting on the presence of the words ‘multiscale’ and ‘multi-scale’ in the project database. For DOE and NIH, we only selected projects that have these phrases directly in the title, while we also searched the abstracts in the case of CORDIS and NSF.

Once we selected the projects, we removed any projects with identical titles, as these are often continuations of the same project in the previous year. Also, we eliminated any

Scientific Domain	Astrophysics	Biology	Energy	Engineering	Environmental	Materials
Distributed multiscale?	few	some	few	unknown	few	few
Acyclic coupling?	some	some	some	most	many	most
Cyclic coupling?	most	most	most	some	many	some
Dominant style of coupling	D	G	G	unknown	D&G	S&D

Table I

ASSESSED CHARACTERISTICS OF THE SIX MULTISCALE SIMULATION DOMAINS, BASED ON THE LITERATURE WE HAVE FOUND. IN THE LAST ROW WE LIST THE MAIN STYLE OF SUBMODEL COUPLING USED IN THESE DISCIPLINES. HERE WE INDICATE DOMAIN-SPECIFIC COUPLING SOLUTIONS WITH A "D", GENERAL-PURPOSE DOMAIN-INDEPENDENT SOLUTIONS WITH A "G", AND COLLECTIONS OF HAND-WRITTEN SCRIPTS WITH AN "S". DUE TO THE COMMERCIAL NATURE OF MANY ENGINEERING MULTISCALE PROJECTS, WE ARE UNSURE ABOUT THE DOMINANT STYLE OF COUPLING OR THE PRESENCE OF DISTRIBUTED MULTISCALE SIMULATIONS IN THAT DOMAIN.




Figure 2. Overview of the spatial and temporal scales in which typical (multiscale) simulations in several scientific domains operate. Each domain is represented as either a colored or a hatched parallelogram.

project that did not describe explicit multiscale modeling or simulation in its abstract. We found over a thousand multiscale simulation and modeling grants, which range from multi-million euro international projects to awards for individual post-doctoral researchers. We provide an overview of these projects by scientific domain in Fig. 3 and by starting year in Fig. 4. The statistics presented here are by no means exhaustive, as we only searched for explicit mentions of multiscale and did not investigate nationally funded projects in the EU, US-based projects funded by other organizations or projects outside both the EU and the US. Our results should therefore be interpreted only as a rough indication of the multiscale community as a whole

and as a lower bound on its size.

In Fig. 3 we find that most multiscale projects reside within the domain of biology and materials. Although there are several multiscale frameworks available in astrophysics, the number of EU projects in this field of research is quite low. This is most likely because international collaboration within theoretical astrophysics tends to focus on more informal international collaborations and national sources of funding.

In Fig. 4 we find that the first multiscale projects emerged in the late 90s, and that the number of these projects in the EU has gradually increased in recent years. The number of multiscale US-based projects peaks in 2009, but has




Figure 3. Overview of multiscale projects by scientific domain. We obtained the data from the EU CORDIS database (cordis.europa.eu), the National Institute of Health (projectreporter.nih.gov), the OSTI database of the Department of Energy (www.osti.gov) and the US National Science Foundation (www.nsf.gov)




Figure 4. Overview of multiscale projects by starting year. We did not find any EU Framework project (multiscale or otherwise) which started in either 1999 or 2007. Additionally, we found no projects in the DOE database which were starting in 2010 or 2011.

diminished in the last few years. This is in part because the DOE database contains no projects starting after 2009 (multiscale or otherwise) and in part because the US Federal Government made a one-time major investment in scientific research in 2009. As most projects often last three years or more, we estimate that there are more than 300 multiscale projects.

B. Coupling frameworks for multiscale simulation

We provide a classification of the coupling frameworks for multiscale simulation in Table II. Here we provide information on whether the tools features a generic implementation, intended to be re-used in other domains, what types of coupling are supported, and whether the tools allow for multiscale simulations run distributed across multiple computational sites. Allowing the distributed execution of multiscale simulations is beneficial, because the subcodes

within each simulation may have heterogeneous resource requirements (e.g., some subcodes may need larger compute resources than others, or require nodes equipped with specialized hardware).

In general, we discern several distinct strategies regarding coupling. Perhaps the most traditional strategy of multiscale coupling is by developing hybrid codes which cover a set of scales within a single simulation code. These *monolithic* codes are generally restricted in their modularity and extensibility. These limitations, combined with the ongoing increase in available compute capacity, have led to a change in coupling strategies and the emergence of more flexible multiscale coupling methods. However, the way these strategies have changed is clearly not uniform across scientific domains.

For example, in the astrophysics and energy domains we observe a focus on reusable domain-specific coupling solutions (e.g., AMUSE and IPS), while in biology and environmental science, the coupling tools developed are generalized in several instances (e.g., MUSCLE and OpenPALM). By generalizing their coupling frameworks, these groups allow their tools to be adopted by researchers in other fields. Researchers in the materials sciences only rarely adopt coupling frameworks, and choose to build their (often acyclically coupled) multiscale simulations either by employing inherent multiscale capabilities within molecular dynamics codes (e.g., by using a “replica exchange” method to model a range of temperatures) or connecting simulations using (often handwritten) pre- and post-processing scripts. In a few instances, however, they do rely on data conversion libraries such as the VOTCA toolkit.

In Table III we provide a basic comparison of important functionalities for multiscale simulations offered by the various coupling frameworks. One of the core functionalities of coupling tools is simplifying access to subcodes by introducing abstractions for the user. All tools that we investigated provide this abstraction on the data level, and a number of tools present standardized interfaces to function calls of individual types of subcodes. The latter functionality is useful because it allows the user to switch solvers in a multiscale simulation with only minor changes in the simulation definition.

Another essential functionality in coupling frameworks is the ability to facilitate data exchange between subcode endpoints and to bootstrap subcodes which are meant to be run concurrently. OpenMI and UAL are mainly intended to improve model access and provide abstractions, and rely on other tools to provide the coupling on a technical level. The GridSpace engine does not inherently allow for bootstrapping concurrently running subcodes, but the user can launch applications that use cyclic coupling by combining GridSpace with a second coupling tool (e.g., MUSCLE). When using coupling frameworks, users are expected to interface to them using common programming languages such

as C, Java or Fortran in most cases. The Swift environment presents a custom-language scripting interface similar to C, while both OpenPALM and GridSpace present a graphical interface to the user.

Overall, most coupling frameworks we have reviewed here provide a wide range of functionalities, and can often be used without supplementary tools to construct multiscale simulations. The MUSCLE, OpenMI, Swift and UAL tools have a more specific scope or are not exclusively aimed at multiscale modelling. As a result, these tools provide a subset of the functionalities required for multiscale simulation.

Using a single heavyweight and domain-specific framework for multiscale simulations is often convenient for the user in the short term, but it comes with several drawbacks on the longer term. First, although it is often straightforward to switch between different solvers within these all-in-one coupling frameworks (sometimes it is as easy as replacing a single line of code), it is often much more difficult to switch from one coupling framework to another. This may be necessary if an existing framework becomes outdated, or if the subcodes within that framework need to be reused outside of the source domain. By constructing and adopting formalizations for defining multiscale coupling patterns (such as MML [54]), we are able to diminish this drawback and improve the portability of multiscale simulations and, for example, allowing them to be more easily moved to a different framework if the existing one becomes obsolete.

Another drawback of using traditional all-in-one approaches is that any new computational improvements in multiscale coupling (such as more powerful data abstractions or improvements in the data exchange performance between subcodes) may have to be applied separately to each framework to be used to full effect, resulting in duplicated integration, or even implementation, efforts. This is a major concern in any large software project, which among other things can be mitigated by strictly enforcing modularity in the framework design (given that the developers of the underlying components use standardized APIs that remain consistent over time).

IV. DISCUSSION AND CONCLUSIONS

We have reviewed a number of multiscale communities and compared them across a range of criteria. The number of multiscale projects has been increasing in recent years so that today there are numerous large multiscale projects in a range of scientific domains. The increase in the number of multiscale projects also implies a growth in the potential benefit that can be gained by developing common and reusable multiscale methods.

The different multiscale communities tend to adopt radically different technical approaches and possess diverse organizational characteristics. Within biology and energy, a considerable fraction of the multiscale projects are bundled in large international initiatives, while the multiscale projects

within astrophysics and materials sciences are often driven by much smaller collaborations. On the technical level, researchers in the astrophysics and energy domains clearly prefer to use domain-specific frameworks to couple their subcodes, while researchers in biology and environmental sciences have a stronger inclination towards general-purpose coupling tools. The numerous projects in the materials sciences adopt yet a different approach, and frequently construct multiscale simulations by connecting codes with hand-written scripts. The vast majority of multiscale simulations are run on single sites, though a small number of projects recently started to explore the use of *distributed multiscale simulations*, where individual subcodes are deployed and run on different computational sites. Considering the heterogeneity in computational requirements of various subcodes, distributed multiscale simulation may be the only way to efficiently run production level multiscale simulations in many cases.

In our analysis of scales simulated by different multiscale computing communities we find a distinct overlap in the scales upon which the simulations in these domains operate. In particular many research domains feature simulations on a length scale of about a meter and a time scale of a few hours. As a result, general-purpose multiscale methods which are geared towards this scale may be particularly suitable for reuse by a wide range of scientific disciplines, and phenomena operating on these scales in one domain may be of non-negligible relevance to others.

It is clear that a uniform strategy for multiscale simulations has yet to emerge, and that different domains have adopted relatively disjoint approaches so far. Nevertheless, multiscale simulations are becoming increasingly widespread to the point where there are at least a few hundred active projects in the EU and the US alone.

It is beyond the scope of this review to pronounce on the benefits of pursuing domain specific approaches versus general purpose approaches for accelerating the progress of multiscale communities. However, based on the findings we presented here, we can clearly conclude that it is high time for such an inter-disciplinary debate to be opened.

ACKNOWLEDGMENTS

We thank our colleagues in the MAPPER consortium and in particular Krzysztof Kurowski, Ilya Saverchenko, Kasia Rycerz, Alfons Hoekstra, Marian Bubak, Werner Dubitzky, Bastien Chopard, David Coster, James Suter and Rupert Nash for their contributions. We also thank Daniel S. Katz, Oliver Henrich and the participants of the Distributed Multiscale Computing workshop 2011 at IEEE e-Science for their valuable input. This work has been supported by the MAPPER EU-FP7 project (grant no. RI-261507).

name	domain of origin	generic implementation?	distributed across sites?	acyclic coupling?	cyclic coupling?	year of last public release
AMUSE [55]	astrophysics	no	yes	yes	yes	2011
FACETS	energy	no	n/a	n/a	yes	2011
FLASH [12]	astrophysics	n/a	n/a	yes	yes	2011
GridSpace [56]	sys. biology	yes	yes	yes	n/a	2011
IPS	energy	no	no	yes	yes	not public
MCT [39]	environment	yes	yes	yes	yes	2009
MUSCLE [57]	sys. biology	yes	yes	n/a	yes	2010
OASIS [41]	environment	no	no	n/a	yes	2011
OpenMI [42]	environment	no	yes	yes	yes	2011
OpenPALM [58]	environment	yes	no	n/a	yes	2011
Pyre [59]	environment	yes	no	yes	yes	2005
Swift [60]	energy	yes	yes	yes	no	2011
SWMF [8]	astrophysics	n/a	no	yes	yes	not public
UAL [23]	energy	yes	yes	yes	yes	not public

Table II

ASSESSED CHARACTERISTICS OF THE TWELVE COUPLING FRAMEWORKS DISCUSSED IN THIS PAPER. ALL THE COUPLING FRAMEWORKS HERE SUPPORT THE SWITCHING AND DYNAMIC USE OF MULTIPLE SUBMODELS IN A MODULAR WAY, AND THE EXECUTION OF PARALLEL MULTISCALE SIMULATIONS WITHIN A SINGLE COMPUTE RESOURCE. WITHIN THE TABLE WE PROVIDE A 'YES' IF THE FRAMEWORK PROVIDES THIS FUNCTIONALITY, 'NO' IF IT CURRENTLY DOES NOT APPEAR TO DO SO, AND 'N/A' IF THE FUNCTIONALITY APPEARS TO BE OUTSIDE OF THE SCOPE OF THE FRAMEWORK ALTOGETHER.

name	abstract subcode function calls	connect subcode endpoints	bootstrap concurrent subcodes	allow sequential workflows	built-in unit conversion?	interface presented to users
AMUSE	yes	yes	yes	yes	yes	Python
FACETS	yes	yes	yes	no	no	C++
FLASH	yes	yes	yes	no	yes	Fortran
GridSpace	no*	yes	no	yes	no	GUI
IPS	yes	yes	yes	yes	no	Python
MCT	no	yes	yes	no	no	Fortran
MUSCLE	no*	yes	yes	no	no	Java
OASIS	no	yes	yes	no	no	Fortran/C
OpenMI	yes	no	no	no	yes	Java/C#
OpenPALM	yes	yes	yes	some	yes	GUI
Pyre	no	yes	yes	no	yes	Python
Swift	no	yes	yes	yes	no	C-like
SWMF	yes	yes	yes	yes	no	Fortran
UAL	no	yes	no	no	no	C/Fortran/Java

Table III

OVERVIEW OF THE FUNCTIONALITIES PROVIDED BY THE COUPLING FRAMEWORKS. TOOLS WHICH PROVIDE A COUPLING ABSTRACTION USING THE MULTISCALE MODELING LANGUAGE [54] ARE MARKED WITH AN ASTERISK IN THE SECOND COLUMN.

REFERENCES

- [1] J. Fish, *Multiscale Methods: Bridging the Scales in Science and Engineering*. Oxford, United Kingdom: Oxford University Press, 2009.
- [2] S. Attinger and P. Koumoutsakos, *Multiscale Modeling and Simulation*. Berlin, Heidelberg, Germany: Springer-Verlag, 2004.
- [3] A. Hoekstra, A. Caiazzo, E. Lorenz, J.-L. Falcone, and

- B. Chopard, *Complex Automata: Multi-scale Modeling with Coupled Cellular Automata*. Springer-Verlag, 2010.
- [4] V. Springel, "The cosmological simulation code GADGET-2," *MNRAS*, vol. 364, pp. 1105–1134, Dec. 2005.
- [5] S. Portegies Zwart, S. McMillan, D. Groen, A. Gualandris, M. Sipior, and W. Vermin, "A parallel gravitational N-body kernel," *New Astronomy*, vol. 13, pp. 285–295, Jul. 2008.
- [6] S. Harfst, A. Gualandris, D. Merritt, and S. Mikkola, "A hybrid N-body code incorporating algorithmic regularization and post-Newtonian forces," *MNRAS*, vol. 389, pp. 2–12, Sep. 2008.
- [7] F. K. Röpke and R. Bruckschen, "Thermonuclear supernovae: a multi-scale astrophysical problem challenging numerical simulations and visualization," *New Journal of Physics*, vol. 10, no. 12, pp. 125 009–+, Dec. 2008.
- [8] G. Tóth, I. V. Sokolov, T. I. Gombosi, D. R. Chesney, C. R. Clauer, D. L. De Zeeuw, K. C. Hansen, K. J. Kane, W. B. Manchester, R. C. Oehmke, K. G. Powell, A. J. Ridley, I. I. Roussev, Q. F. Stout, O. Volberg, R. A. Wolf, S. Sazykin, A. Chan, B. Yu, and J. Kóta, "Space Weather Modeling Framework: A new tool for the space science community," *Journal of Geophysical Research (Space Physics)*, vol. 110, no. A9, p. A12226, dec 2005.
- [9] S. Portegies Zwart, S. McMillan, I. Pelupessy, and A. van Elteren, "Multi-physics Simulations using a Hierarchical Interchangeable Software Interface," in *Advances in Computational Astrophysics: Methods, Tools, and Outcome*, ser. Astronomical Society of the Pacific Conference Series, R. Capuzzo-Dolcetta, M. Limongi, and A. Tornambè, Eds., vol. 453, Jul. 2012, p. 317.
- [10] S. Portegies Zwart, S. McMillan, S. Harfst, D. Groen, M. Fujii, B. Ó Nualláin, E. Glebbeek, D. Hoggie, J. Lombardi, P. Hut, V. Angelou, S. Banerjee, H. Belkus, T. Fragos, J. Fregeau, E. Gaburov, K. Izzard, M. Juric, S. Justham, A. Sottoriva, P. Teuben, J. van Bever, O. Yaron, and M. Zemp, "A multiphysics and multiscale software environment for modeling astrophysical systems," *New Astronomy*, vol. 14, no. 4, pp. 369 – 378, 2009.
- [11] B. Fryxell, K. Olson, P. Ricker, F. X. Timmes, M. Zingale, D. Q. Lamb, P. MacNeice, R. Rosner, J. W. Truran, and H. Tufo, "Flash: An adaptive mesh hydrodynamics code for modeling astrophysical thermonuclear flashes," *The Astrophysical Journal Supplement Series*, vol. 131, no. 1, p. 273, 2000.
- [12] "Flash," 2011. [Online]. Available: [FLASHmultiscaleframework:http://flash.uchicago.edu/website/home/](http://flash.uchicago.edu/website/home/FLASHmultiscaleframework)
- [13] D. Noble, "Modeling the heart—from genes to cells to the whole organ," *Science*, vol. 295, no. 5560, pp. 1678–1682, 2002.
- [14] A. Finkelstein, J. Hetherington, L. Li, O. Margoninski, P. Saffrey, R. Seymour, and A. Warner, "Computational challenges of systems biology," *Computer*, vol. 37, pp. 26–33, May 2004.
- [15] J. Hetherington, I. D. L. Bogle, P. Saffrey, O. Margoninski, L. Li, M. Varela Rey, S. Yamaji, S. Baigent, J. Ashmore, K. Page, R. M. Seymour, A. Finkelstein, and A. Warner, "Addressing the challenges of multiscale model management in systems biology," *Computers and Chemical Engineering*, vol. 31, no. 8, pp. 962 – 979, 2007, jce:title;7th World Congress of Chemical Engineering/cce:title;.
- [16] J. Southern, J. Pitt-Francis, J. Whiteley, D. Stokeley, H. Kobashi, R. Nobes, Y. Kadooka, and D. Gavaghan, "Multi-scale computational modelling in biology and physiology," *Progress in Biophysics and Molecular Biology*, vol. 96, no. 1-3, pp. 60 – 89, 2008, cardiovascular Physiome.
- [17] J. Pitt-Francis, P. Pathmanathan, M. O. Bernabeu, R. Bordas, J. Cooper, A. G. Fletcher, G. R. Mirams, P. Murray, J. M. Osborne, A. Walter, S. J. Chapman, A. Garny, I. M. M. van Leeuwen, P. K. Maini, B. Rodriguez, S. L. Waters, J. P. Whiteley, H. M. Byrne, and D. J. Gavaghan, "Chaste: A test-driven approach to software development for biological modelling," *Computer Physics Communications*, vol. 180, no. 12, pp. 2452 – 2471, 2009, jce:title;40 YEARS OF CPC: A celebratory issue focused on quality software for high performance, grid and novel computing architectures/cce:title;.
- [18] S. Schnell, R. Grima, and P. K. Maini, "Multiscale modeling in biology," *American Scientist*, vol. 95, pp. 134–142, 2007.
- [19] P. M. A. Sloot and A. G. Hoekstra, "Multi-scale modelling in computational biomedicine," *Briefings in Bioinformatics*, vol. 11, no. 1, pp. 142–152, 2010.
- [20] J. O. Dada and P. Mendes, "Multi-scale modelling and simulation in systems biology," *Integr. Biol.*, vol. 3, pp. 86–96, 2011. [Online]. Available: <http://dx.doi.org/10.1039/C0IB00075B>
- [21] P. J. Hunter, E. J. Crampin, and P. M. F. Nielsen, "Bioinformatics, multiscale modeling and the iups physiome project," *Briefings in Bioinformatics*, vol. 9, no. 4, pp. 333–343, 2008.
- [22] D. J. Hill, "Nuclear energy for the future," *Nature Materials*, vol. 7, pp. 680–682, Sep. 2008.
- [23] G. Manduchi, F. Iannone, F. Imbeaux, G. Huysmans, J. Lister, B. Guillerminet, P. Strand, L.-G. Eriksson, and M. Romanelli, "A universal access layer for the integrated tokamak modelling task force," *Fusion Engineering and Design*, vol. 83, no. 2-3, pp. 462 – 466, 2008.
- [24] M. Wilde, I. Foster, K. Iskra, P. Beckman, Z. Zhang, A. Espinosa, M. Hategan, B. Clifford, and I. Raicu, "Parallel scripting for applications at the petascale and beyond," *Computer*, vol. 42, pp. 50–60, November 2009. [Online]. Available: <http://dx.doi.org/10.1109/MC.2009.365>
- [25] "FLOMANIA: flow physics modelling, an integrated approach." 2011. [Online]. Available: <http://cfd.mace.manchester.ac.uk/flomania/>
- [26] "DESIDER: detached eddy simulations for industrial aerodynamics." 2011. [Online]. Available: <http://cfd.mace.manchester.ac.uk/desider/>

- [27] D. F. Martin, P. Colella, M. Anghel, and F. J. Alexander, "Adaptive mesh refinement for multiscale nonequilibrium physics," *Computing in Science and Engineering*, vol. 7, pp. 24–31, 2005.
- [28] M.-O. Coppens and G. Wang, *Optimal Design of Hierarchically Structured Porous Catalysts*. Wiley-VCH Verlag GmbH and Co. KGaA, 2009, pp. 25–58.
- [29] D. Vlachos and S. Caratzoulas, "The roles of catalysis and reaction engineering in overcoming the energy and the environment crisis," *Chemical Engineering Science*, vol. 65, no. 1, pp. 18 – 29, 2010, [jce:title20th International Symposium in Chemical Reaction EngineeringGreen Chemical Reaction Engineering for a Sustainable Futurejce:title2](#).
- [30] E. Cueto, M. Laso, and F. Chinesta, "Meshless stochastic simulation of micro-macrokinetic theory models," *International Journal of Multiscale Computational Engineering*, vol. 9, no. 1, pp. 1–16, 2011.
- [31] V. Symeonidis, G. E. Karniadakis, and B. Caswell, "A seamless approach to multiscale complex fluid simulation," *Computing in Science and Engineering*, vol. 7, pp. 39–46, 2005.
- [32] H. S. Wijesinghe and N. G. Hadjiconstantinou, "Discussion of hybrid atomistic-continuum methods for multiscale hydrodynamics," *International Journal for Multiscale Computational Engineering*, vol. 2, no. 2, 2004.
- [33] P. van Thang, B. Chopard, L. Lefèvre, D. A. Ondo, and E. Mendes, "Study of the 1D lattice Boltzmann shallow water equation and its coupling to build a canal network," *Journal of Computational Physics*, vol. 229, pp. 7373–7400, Sep. 2010.
- [34] B. Plale, D. Gannon, J. Brotzge, K. Droegemeier, J. Kurose, D. McLaughlin, R. Wilhelmson, S. Graves, M. Ramamurthy, R. D. Clark, S. Yalda, D. A. Reed, E. Joseph, and V. Chandrasekar, "Casa and lead: Adaptive cyberinfrastructure for real-time multiscale weather forecasting," *IEEE Computer*, vol. 39, pp. 56–64, 2006.
- [35] R. M. Winglee, E. Harnett, and A. Kidder, "Relative timing of substorm processes as derived from multifluid/multiscale simulations: Internally driven substorms," *Journal of Geophysical Research (Space Physics)*, vol. 114, no. A13, p. A09213, sep 2009.
- [36] "Community climate system model," 2011. [Online]. Available: <http://www.scidac.gov/CCSM/>
- [37] C. An and Y. Cai, "The effect of beach slope on the tsunami run-up induced by thrust fault earthquakes," *Procedia Computer Science*, vol. 1, no. 1, pp. 645 – 654, 2010, iCCS 2010.
- [38] R. Klein, S. Vater, E. Paeschke, and D. Ruprecht, *Multiple scales methods in meteorology*. Springer Verlag, 2011, pp. 127–196.
- [39] J. Larson, R. Jacob, and E. Ong, "The model coupling toolkit: A new Fortran90 toolkit for building multiphysics parallel coupled models," *International Journal of High Performance Computing Applications*, vol. 19, no. 3, pp. 277–292, Fall 2005.
- [40] E. Tan, E. Choi, P. Thoutireddy, M. Gurnis, and M. Aivazis, "Geoframework: Coupling multiple models of mantle convection within a computational framework," *Geochem. Geophys. Geosyst.*, vol. 7, no. Q06001, 2006.
- [41] R. Redler, S. Valcke, and H. Ritzdorf, "OASIS4 - a coupling software for next generation earth system modelling," *Geoscientific Model Development*, vol. 3, pp. 87–104, Jan. 2010.
- [42] J. B. Gregersen, P. J. A. Gijssbers, and S. J. P. Westen, "OpenMI: Open modelling interface," *Journal of Hydroinformatics*, vol. 9, no. 3, pp. 175–191, 2007.
- [43] "RADINTERFACES," 2011. [Online]. Available: <http://www.materials.imdea.org/Research/Projects/RADINTERFACE/tabid/3>
- [44] "M3-2s," 2011. [Online]. Available: <http://www.m3-2s.bham.ac.uk/>
- [45] J. R. Kermode, T. Albaret, D. Sherman, N. Bernstein, P. Gumbsch, M. C. Payne, G. Csányi, and A. de Vita, "Low-speed fracture instabilities in a brittle crystal," *Nature*, vol. 455, pp. 1224–1227, Oct. 2008.
- [46] E. Weinan, B. Engquist, X. Li, W. Ren, and E. Vanden-Eijnden, "Heterogeneous Multiscale Methods: A Review," *Communications in Computational Physics*, vol. 2, no. 3, pp. 367–450, Jun. 2007.
- [47] D. Groen, J. Suter, and P. V. Coveney, "Modelling distributed multiscale simulation performance: An application to nanocomposites," in *Seventh IEEE international conference on e-Science and Grid computing: Stockholm, Sweden*. Piscataway, NJ: IEEE Computer Society, December 2011, pp. 105–111.
- [48] A. Donev, J. B. Bell, A. L. Garcia, and B. J. Alder, "A hybrid particle-continuum method for hydrodynamics of complex fluids," *Multiscale Model. Simul.*, vol. 8, pp. 871–911, 2010.
- [49] R. Delgado-Buscalioni and P. V. Coveney, "Continuum-particle hybrid coupling for mass, momentum, and energy transfers in unsteady fluid flow," *Phys. Rev. E*, vol. 67, p. 046704, Apr 2003.
- [50] M. Bernaschi, S. Melchionna, S. Succi, M. Fyta, E. Kaxiras, and J. K. Sircar, "MUPHY: A parallel MULTI PHYsics/scale code for high performance bio-fluidic simulations," *Computer Physics Communications*, vol. 180, pp. 1495–1502, Sep. 2009.
- [51] Y.-B. Peng and J. Li, "Multiscale analysis of stochastic fluctuations of dynamic yield of magnetorheological fluids," *International Journal for Multiscale Computational Engineering*, vol. 9, no. 2, pp. 175–191, 2011.
- [52] "International journal for multiscale computational engineering," 2012. [Online]. Available: <http://www.begellhouse.com/journals/61fd1b191cf7e96f.html>
- [53] "Journal of multiscale modelling," 2012. [Online]. Available: <http://www.worldscinet.com/jmm/>

- [54] J. Borgdorff, Falcone, J.-L., E. Lorenz, B. Chopard, and A. G. Hoekstra, "A principled approach to distributed multiscale computing, from formalization to execution," in *Seventh IEEE international conference on e-Science and Grid computing: Stockholm, Sweden (in press)*. Piscataway, NJ: IEEE Computer Society, December 2011.
- [55] "Amuse, the astrophysical multipurpose software environment," 2011. [Online]. Available: <http://www.amusecode.org>
- [56] "Distributed computing environments - gridspace technology," 2011. [Online]. Available: <http://dice.cyfronet.pl/gridspace/>
- [57] J. Hegewald, M. Krafczyk, J. Tölke, A. Hoekstra, and B. Chopard, "An agent-based coupling platform for complex automata," in *Proceedings of the 8th international conference on Computational Science, Part II*, ser. ICCS '08. Berlin, Heidelberg: Springer-Verlag, 2008, pp. 227–233.
- [58] "OpenPALM," 2012. [Online]. Available: http://www.cerfacs.fr/globc/PALM_WEB/
- [59] "The pyre framework," 2012. [Online]. Available: <http://www.cacr.caltech.edu/projects/pyre/>
- [60] M. Wilde, M. Hategan, J. M. Wozniak, B. Clifford, D. S. Katz, and I. Foster, "Swift: A language for distributed parallel scripting," *Parallel Computing*, vol. 39, no. 9, pp. 633–652, September 2011.